



PCI-SIG ENGINEERING CHANGE NOTICE

| | |
|---------------------------|--|
| TITLE: | Enhanced PCIe Precision Time Measurement (ePTM) |
| DATE: | Introduced: 2 May 2018 Updated: 8 January 2019 Final Approval: 10 January 2019 |
| AFFECTED DOCUMENT: | PCI Express 4.0 Base Specification |
| SPONSOR: | Intel Corporation |

Part I

1. Summary of the Functional Changes

ePTM is an improvement on the existing Precision Time Measurement capability that provides improved detection and handling of error cases. ePTM quickly identifies and resolves errors that may cause clocks to become desynchronized.

2. Benefits as a Result of the Changes

A design implementing ePTM will benefit from improved detection and handling of error cases that would otherwise cause clocks to become desynchronized.

3. Assessment of the Impact

ePTM only affects designs that utilize PTM. See 4 and 5 for additional details.

4. Analysis of the Hardware Implications

ePTM is fully backwards compatible with existing PTM implementations. PTM and ePTM implementations are interoperable. PTM enabled implementations receive the improved error detection of ePTM when connected to a ePTM implementation.

5. Analysis of the Software Implications

ePTM uses the same software interface as PTM. Existing PTM software should be compatible with new ePTM implementations. No software changes should be necessary to enable ePTM.

The only change visible to software is the addition of the ePTM Capable bit to the PTM Extended Capability.

6. Analysis of the C&I Test Implications

This ECN allocates previously reserved bits and requires the creation of new tests for the new register fields. ePTM Asserts additional operational requirements on PTM Requesters and PTM Responders. These additional requirements are a superset of the existing PTM requirements. New C & I tests would be required if it is desirable to extend C & I coverage to explicitly evaluate these new requirements.

Make the following changes to Section 6.22.1:

6.22.1 Introduction

Precision Time Measurement (PTM) enables precise coordination of events across multiple components with independent local time clocks. Ordinarily, such precise coordination would be difficult given that individual time clocks have differing notions of the value and rate of change of time. To work around this limitation, PTM enables components to calculate the relationship between their local times and a shared PTM Master Time: an independent time domain associated with a PTM Root.

Enhanced Precision Time Management (ePTM) places additional requirements on PTM Devices. Support for ePTM is indicated by the ePTM Capable bit.

PTM defines the following:

- PTM Requester - A Function capable of using PTM as a consumer associated with an Endpoint or an Upstream Port.
- PTM Responder - A Function capable of using PTM to supply PTM Master Time associated with a Port or an RCRB.
- Time Source - A local clock associated with a PTM Responder.
- PTM Root - The source of PTM Master Time for a PTM Hierarchy. A PTM Root must also be a Time Source and is typically also a PTM Responder.

Each PTM Root supplies a single PTM Master Time ~~for a~~ to all of the PTM Hierarchy: a set of PTM Requesters associated with a single PTM Root.

...

Add the following Implementation Note to the end of Section 6.22.1:

IMPLEMENTATION NOTE: PTM and Retimers

PCIe Retimers can impact PTM accuracy by introducing asymmetric link delays. Retimers designed to maintain symmetric link delays will enable the best PTM accuracy. The larger and more variable the asymmetry, the greater the impact to PTM. Consult the manufacturer's documentation to determine the suitability of a Retimer implementation for use with PTM.

Make the following changes to Section 6.22.3.1:

6.22.3.1 PTM Requester Role

- Support for the PTM Requester role is indicated by setting the PTM Requester Capable bit in the PTM Capability register.
- PTM Requesters are permitted to request PTM Master Time only when PTM is enabled. The mechanism for directing a PTM Requester to issue such a request is implementation specific.

- Upstream Ports obtain PTM Master Time via PTM dialogs as described in section 2.2.8.10.
- The mechanism by which RCiEPs request PTM Master Time is implementation specific.
- Once having issued a PTM Request Message, the Upstream Port must not issue another PTM Request Message prior to the receipt of a PTM Response Message, PTM ResponseD Message, Reset, or the passage of 100 μ s without a corresponding PTM Message from the Downstream Port.
- Upon receiving a PTM Response, the Upstream Port must wait at least 1 μ s before issuing another PTM Request Message.
- For Multi-Function Devices (MFDs) containing multiple PTM Requesters, the Upstream Port associated with that MFD must issue a single PTM dialog during each PTM context refresh. PTM Requesters within the MFD maintain their individual PTM contexts using this one, Device-wide PTM dialog. The mechanism for refreshing multiple PTM contexts from one PTM dialog is implementation specific.
- It is strongly recommended that an Upstream Port invalidate its internal PTM context when any of the following occur. If ePTM is supported, then an Upstream Port must invalidate its internal PTM context when any of the following occur:
 - A PTM Request is replayed.
 - A duplicate PTM ResponseD TLP is received.
 - The relationship between PTM Master Time and the Upstream Port's local time changes, as determined by implementation specific criteria. For example, this may occur as a result of a transition to a non-D0 state or due to accumulated PPM drift.

These events are grouped under the label "Local Time Invalidation Event" in Figure 6-26.

- If ePTM is supported, an Upstream Port, upon replaying a PTM TLP, must invalidate it's PTM context until two successive PTM dialogs have been completed successfully and without replays.

...

Add the following Implementation Note to the end of Section 6.22.3.1:

IMPLEMENTATION NOTE: PTM Invalidation on the Reception of duplicate TLPs

Duplicate TLPs are detected and discarded in the Data Link Layer, whereas PTM messages are identified in the Transaction Layer. In some implementations it may be difficult or excessively complicated to distinguish a duplicate PTM TLP from other duplicate TLPs.

Because Upstream Ports are permitted to invalidate their internal PTM context for implementation-specific criteria, a PTM Requester is allowed to invalidate its internal PTM context upon the reception of any duplicate TLP in addition to any duplicate PTM

TLP. Similarly, if ePTM is supported, then a PTM Responder is allowed to invalidate its historical timestamps (t2-t3) upon the reception of any duplicate TLP.

Make the following changes to Section 6.22.3.2:

6.22.3.2 PTM Responder Role

- Support for the PTM Responder role is indicated by setting the PTM Responder Capable bit in the PTM Capability register.
- Switches and Root Complexes are permitted to implement the PTM Responder Role.
 - A PTM capable Switch, when enabled for PTM by setting the PTM Enable bit in the PTM Control register associated with the Switch Upstream Port, must respond to all PTM Request Messages received at any of its Downstream Ports.
 - The mechanism by which Root Complexes communicate PTM Master Time to RCiEPs is implementation specific.
- PTM Responders must populate PTM ResponseD Messages as follows (refer to Figure 6-24 and the accompanying implementation note):
 - The PTM Master Time field is a 64-bit value containing the value of PTM Master Time at the receipt of the PTM Request Message for the current PTM Dialog. In Figure 6-24 for the 2nd PTM dialog, this is the PTM Master Time at time t2'.
 - The Propagation Delay field is a 32-bit value containing the interval between the receipt of the PTM Request Message and the transmission of the PTM Response Message for the previous PTM dialog. In Figure 6-24 , for the 2nd PTM dialog, this is the time interval between t2 and t3 captured during the 1st PTM dialog.
 - The unit of measurement for both fields is one ns.
 - A PTM Responder with multiple downstream ports must populate all PTM ResponseD Messages with values from a single PTM Root across all its PTM Ports Downstream ports.
- Switch Downstream Ports and Root Ports acting as PTM Responders must respond to each PTM Request Message received at their Downstream Ports with either PTM Response or PTM ResponseD according to the following rules:
 - A PTM Responder must not send a PTM Response or PTM ResponseD Message without first receiving a PTM Request Message.
 - Upon receipt of a PTM Request Message, a PTM Responder must attempt to issue a PTM Response or PTM ResponseD Message within 10 μ s.
 - A PTM Responder must issue PTM Response when the Downstream Port does not have valid historical timestamps (t3 - t2) with which to fulfill a PTM Request Message.
 - If ePTM is supported, a PTM Responder must invalidate its historical timestamps (t3 - t2) immediately upon replaying any PTM Response or PTM ResponseD. A PTM Responder must invalidate its historical timestamps (t3 - t2) after receiving any duplicate PTM Request.

- A PTM Responder must issue PTM ResponseD when it has stored copies of the values required to populate the PTM ResponseD Message: historical timestamps (~~t3-t2~~) - t2) and the PTM Master Time at the receipt of the most recent PTM Request Message (time t2’).
- A PTM Responder is permitted to issue PTM Response when it has stored copies of the historical timestamps (t3 - t2) but must request the PTM Master Time from elsewhere. In this case, it is permitted to issue PTM Response messages in response to PTM Request Messages while it retrieves the PTM Master Time if that retrieval is expected to take more than 10 µs.
- The perceived granularity of the historical timestamps and PTM Master Time values transmitted by a PTM Responder must not exceed that reported in the Local Clock Granularity field of the PTM Capability register.

...

Make the following changes to Table 7-200:

Editor’s Note: Update Figure 7-246, to add new bit 3, removing it from the RsvdP field

Table 7-200 PTM Capability Register

| Bit Location | Register Description | Attributes |
|--------------|---|---------------|
| 0 | <p>PTM Requester Capable - <u>Indicates the Function implements the PTM Requester role (see Section 6.22.3.1).</u></p> <p>Endpoints and RCiEPs are permitted to, and Switches supporting PTM must, Set this bit to 1b to indicate <u>that</u> they implement the PTM Requester role (see Section 6.22.3.1).</p> <p><u>Switch Upstream Ports must Set this bit if the Switch contains one or more of the following:</u></p> <ul style="list-style-type: none"> • <u>A Downstream Port that implements the PTM Responder role.</u> • <u>An additional Function that implements the PTM Requester role.</u> <p><u>If a Device contains multiple Upstream Port Functions, The value of this bit must be consistent across all such Functions.</u></p> | HwInit |
| 1 | <p>PTM Responder Capable - Root Ports and RCRB’s are permitted to, and Switches supporting PTM must, Set this bit to 1b to indicate they implement the PTM Responder role (see Section 6.22.3.2).</p> <p>If PTM Root Capable is Set; <u>then</u> this bit must be Set to 1b.</p> | HwInit |
| 2 | <p>PTM Root Capable - Root Ports, RCRBs and Switches are permitted to Set this bit to 1b, if they implement a PTM Time Source Role and are capable of serving as the PTM Root. <u>being a source of PTM Master Time (see Section 6.22.1).</u></p> <p><u>All other Functions must hardwire this bit to 0b.</u></p> | HwInit |
| <u>3</u> | <p><u>ePTM Capable - If Set, indicates that this device supports Enhanced Precision Time Management (ePTM).</u></p> <p><u>It is strongly recommended that this bit be Set in all PTM Devices.</u></p> | <u>HwInit</u> |

Table 7-200 PTM Capability Register

| Bit Location | Register Description | Attributes |
|--------------|---|----------------|
| 15:8 | <p><u>Editor's Note: Formatting Changes (Bold title & tableized enumerated values)</u></p> <p>Local Clock Granularity - Encodings are:</p> <p>0000 0000b— Time Source does not implement a local clock. It simply propagates timing information obtained from further Upstream in the PTM Hierarchy when responding to PTM Request messages.</p> <p>0000 0001b – Indicates the period of this Time Source's local clock in ns. 1111 1110b÷</p> <p>1111 1111b÷ Indicates the period of this Time Source's local clock is greater than 254 ns.</p> <p>If the PTM Root Select bit is Set, this local clock is used to provide PTM Master Time. Otherwise, the Time Source uses this local clock to locally track PTM Master Time received from further Upstream within a PTM Hierarchy.</p> <p>This field is reserved for Functions that do not implement the PTM Time Source role RsvdP if the PTM Root Capable bit is 0b.</p> | HwInit / RsvdP |